

JUN-25-2007 MON 02:10 PM TROP, PRUNER & HU, PC

FAX NO. 7134688883

RECEIVED
CENTRAL FAX CENTER

P. 01/23

JUN 25 2007

TROP, PRUNER & HU, P.C.

INTELLECTUAL PROPERTY LAW ATTORNEYS

1616 S. Voss Rd., Suite 750
Houston, Texas 77057-3621
Bus: (713) 468-8880
Fax: (713) 468-8883

Fax

To: Examiner Robert M. Timblin **From:** Dan C. Hu

Company: U.S. Patent Office **Date:** June 25, 2007

Fax: (571) 273-8300 **Pages:** 23 (including cover sheet)

Serial No.: 10/706,656 **Our Ref:** 11166 (NCR.0114US)

Urgent **For Review** **Please Comment** **Please Reply** **Confirm Receipt**

RECEIVED IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Attorney Docket No.: 11166 (NCR.0114US)
Date: June 25, 2007

DCH/gky

Applicant(s): HONG GUI ET AL.

Serial No.: 10/706,656

Filing Date: November 12, 2003

Title: MAINTAINING VIEWS OF CUBE-BASED OPERATIONS IN A DATABASE SYSTEM

1. Appeal Brief Transmittal (in duplicate); and
2. Appeal Brief.

● **Notice:** This information is intended to be for the use of the individual or entity named on this transmittal sheet. If you are not the intended recipient, be aware that any disclosure, copying, distribution, or use of the contents of this faxed information is prohibited. If you have received this facsimile in error, please notify the sender by telephone immediately so that arrangements can be made for the retrieval of the original document at no cost to you.

JUN-25-2007 MON 02:10 PM TROP, PRUNER & HU, PC

FAX NO. 7134688883

RECEIVED P. 02/23
CENTRAL FAX CENTER

JUN 25 2007

TRANSMITTAL OF APPEAL BRIEF (Large Entity)

Docket No.
11166 (NCR.0114US)

In Re Application Of: Hong Gui et al.

Application No.	Filing Date	Examiner	Customer No.	Group Art Unit	Confirmation No.
10/706,656	11-12-2003	Robert M. Timblin	26890	2167	8622

Invention: Maintaining Views of Cube-Based Operations in a Database System

COMMISSIONER FOR PATENTS:

Transmitted herewith is the Appeal Brief in this application, with respect to the Notice of Appeal filed on:
April 25, 2007.

The fee for filing this Appeal Brief is: \$500.00

- A check in the amount of the fee is enclosed.
- The Director has already been authorized to charge fees in this application to a Deposit Account.
- The Director is hereby authorized to charge any fees which may be required, or credit any overpayment to Deposit Account No. 14-0225 (11166). I have enclosed a duplicate copy of this sheet.
- Payment by credit card. Form PTO-2038 is attached.

WARNING: Information on this form may become public. Credit card information should not be included on this form. Provide credit card information and authorization on PTO-2038.

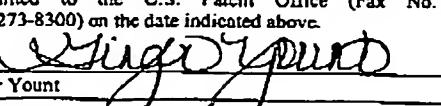


Signature

Dated: June 25, 2007

Dan C. Hu
Registration No. 40,025
TROP, PRUNER & HU, P.C.
1616 South Voss Road, Suite 750
Houston, TX 77057-2631
Telephone: (713) 468-8880, ext. 304
Facsimile: (713) 468-8883

06/26/2007 TL0111 08888825 140225 10786656
cc: 01 FC:1402 500.00 DA

Date of Deposit:	June 25, 2007
I hereby certify that this correspondence is being facsimile transmitted to the U.S. Patent Office (Fax No. (571) 273-8300) on the date indicated above.	
 Ginger Yount	

RECEIVED
CENTRAL FAX CENTER
JUN 25 2007

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants:	Hong Gui et al.	§ Art Unit:	2167
Serial No.:	10/706,656	§	
Filed:	November 12, 2003	§ Examiner:	Tarik Koc
For:	Maintaining Views of Cube-Based Operations in a Database System	§ Atty. Dkt. No.:	11166 (NCR.0114US)
		§	

Mail Stop Appeal Brief-Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF PURSUANT TO 37 C.F.R § 41.37

Sir:

The final rejection of claims 1-28 is hereby appealed.

I. REAL PARTY IN INTEREST

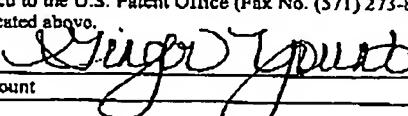
The real party in interest is NCR Corporation.

II. RELATED APPEALS AND INTERFERENCES

None.

III. STATUS OF THE CLAIMS

Claims 1-28 have been finally rejected and are the subject of this appeal. A Reply to Final Office Action was filed on February 23, 2007. However, no claim amendment was made in the Reply to Final Office Action dated February 23, 2007.

Date of Deposit:	<u>June 25, 2007</u>
I hereby certify that this correspondence is being facsimile transmitted to the U.S. Patent Office (Fax No. (571) 273-8300) on the date indicated above.	
	
Ginger Yount	

RECEIVED
CENTRAL FAX CENTER

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

JUN 25 2007

IV. STATUS OF AMENDMENTS

No amendment after final rejection has been submitted.

V. SUMMARY OF THE CLAIMED SUBJECT MATTER

The following provides a concise explanation of the subject matter defined in each of the independent claims involved in the appeal, referring to the specification by page and line number and to the drawings by reference characters, as required by 37 C.F.R. § 41.37(c)(1)(v). Each element of the claims is identified by a corresponding reference to the specification and drawings where applicable. Note that the citation to passages in the specification and drawings for each claim element does not imply that the limitations from the specification and drawings should be read into the corresponding claim element.

Independent claim 1 recites a database system comprising:

a storage (Fig. 1:104) to store a view (Fig. 1:105) containing results of a cube-based operation (Spec., ¶ [016]) on at least one base table, the view containing a first result set for a group-by on a first grouping set, and a second result set for a group-by on a second grouping set (Spec., ¶ [018], [027]); and

a controller (Fig. 1:112, 100), in response to a change to the at least one base table, to:

update the first result set by computing a change to the first result set based on a change in the at least one base table (Spec., ¶ [035], [047]-[050]); and

update the second result set by computing a change to the second result set based on the change to the first result set (Spec., ¶ [035], [047]-[050]).

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

Independent claim 13 recites a method for use in a database system, comprising:

storing a view (Fig. 1:105) containing results of a cube-based operation (Spec., ¶ [016]) on at least one base table, the view containing result sets for group-bys on respective grouping sets (Spec., ¶¶ [018], [027]);

updating a first result set by computing a change to the first result set based on a change in the at least one base table (Spec., ¶¶ [035], [047]-[050]); and

updating a second result set by computing a change to the second result set based on the change to the first result set (Spec., ¶¶ [035], [047]-[050]).

Independent claim 22 recites an article comprising at least one storage medium containing instructions that when executed cause a database system to:

store a view (Fig. 1:105) containing results of a cube-based operation on at least one base table, the view containing result sets for group-bys on respective grouping sets (Spec., ¶¶ [018], [027]);

update a first result set by computing a change to the first result set based on a change in the at least one base table (Spec., ¶¶ [035], [047]-[050]); and

update a second result set by computing a change to the second result set based on the change to the first result set (Spec., ¶¶ [035], [047]-[050]).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

- A. **Claims 1-5, 13-16, And 22-24 Rejected Under 35 U.S.C. § 102 Over U.S. Patent Application Publication No. 2003/0093407 (Cochrane).**
- B. **Claims 6-12, 17-21, And 25-28 Rejected Under 35 U.S.C. § 103 Over Cochrane In View Of U.S. Patent No. 7,035,843 (Bellamkonda).**

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

VII. ARGUMENT

The claims do not stand or fall together. Instead, Appellant presents separate arguments for various independent and dependent claims. Each of these arguments is separately argued below and presented with separate headings and sub-headings as required by 37 C.F.R. § 41.37(c)(1)(vii).

A. Claims 1-5, 13-16, And 22-24 Rejected Under 35 U.S.C. § 102 Over U.S. Patent Application Publication No. 2003/0093407 (Cochrane).

1. Claims 1, 2, 4, 5, 13, 14, 22, and 23.

Independent claim 1 was rejected as being anticipated by Cochrane. Appellant respectfully submits that Cochrane does not disclose each and every element of claim 1, and therefore, does not anticipate claim 1.

Claim 1 recites a database system comprising:

- a storage to store a view containing results of a cube-based operation on at least one base table, the view containing a first result set for a group-by on a first grouping set, and a second result set for a group-by on a second grouping set; and
- a controller, in response to a change to the at least one base table, to:
 - update the first result set by computing a change to the first result set based on a change in the at least one base table; and
 - update the second result set by computing a change to the second result set based on the change to the first result set.

With respect to the controller element of claim 1, the Examiner cited ¶¶ [0035] and [0038] of Cochrane as disclosing the two update tasks performed by the controller of claim 1. 12/27/2006 Office Action at 3-4. Appellant respectfully disagrees with this assertion by the Examiner.

Paragraph [0038] of Cochrane refers to materialized views based on a group-by clause expression specifying CUBE(), ROLLUP(), and GROUPING SETS(). Paragraph [0035] of

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

Cochrane indicates that maintenance of materialized views is incremental in nature to avoid full recomputation of the queries of the materialized views. The same paragraph mentions that a self-maintainable materialized view can be maintained incrementally by processing only the set of updated rows without reference to the underlying tables of the materialized view query.

The incremental update of a materialized view is explained more fully on pages 6 and 7 of Cochrane. On page 6 of Cochrane, ¶ [0093] refers to a propagation phase in which a raw delta (representing a change) as a result of an insert/delete/update operation on the base table is computed. Paragraph [0094] of Cochrane then states that the raw delta stream is aggregated according to the grouping definition of the summary table (view). Significantly, the aggregated delta stream includes multiple, distinct grouping combinations. Cochrane, ¶ [0098]. This means that the delta stream is aggregated according to the grouping specification of the summary table (view). Cochrane, ¶ [0108].

It is of *paramount importance* for that patent application that this step results in a data stream *holding multiple grouping combinations* in the presence of a complex grouping expression of the summary table. For example, if the summary table is defined using 'CUBE()', then this aggregation step yields in a complete *delta cube* with 'higher' aggregate values for the original delta changes.

Id. (emphasis added).

Thus, the delta aggregation of Cochrane produces a "delta cube" according to the complex group-by expression of the view definition. Cochrane, ¶ [0110]. After aggregation, this delta cube "is paired with the current content of the summary table [view] using a left outer-join over the grouping and grouping function columns" Cochrane, ¶ [0114].

In other words, it is clear that what Cochrane contemplates is that after the changes to a base table have been identified (referred to as the delta), that delta is aggregated according to the grouping sets specified in the complex group-by expression to produce a delta cube, such that a

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

one-to-one correspondence between the content of the delta cube and the rows in the view to be updated exists.

This type of incremental update of a view, as performed by Cochrane, is completely different from what is recited in claim 1, where the controller, in response to a change to the at least one base table, updates the first result set by computing a change to the first result set based on a change in the at least one base table. However, the second result set is updated by the controller by computing a change to the second result set *based on the change to the first result set*. In contrast, in Cochrane, updates of all entries of the view that are affected by a base table change are based on the change to the base table, by applying aggregation of the changes to the base table to produce a delta cube according to the grouping sets specified in the complex group-by expression. Clearly, updating the different portions of the view in Cochrane based on the delta cube does not constitute updating one result set based on a change in another result set (where both result sets correspond to respective grouping sets).

Rather than refer to the specific and detailed explanation (on pages 6 and 7) of how materialized views are incrementally maintained in Cochrane, the Examiner in the Office Action made the following observation with respect to ¶ [0035]:

Paragraph 0035 discloses a method in which changes (the equivalent of updates) can be maintained incrementally by processing the updated rows and visiting the underlying tables [of the materialized view query] to recompute some portions of the materialized view query. Therefore, an update affecting a second result set would be calculated by visiting the already materialized view query to obtain pre-computed rows of data (the first result set) in order to compute an update to a second result set, meeting the limitation computing a change to the second result set based on the change to the first result set.

The statement above that “an update affecting a second result set would be calculated by visiting the already materialized view query to obtain pre-computed rows of data (the first result set) in order to compute an update to a second result set,” is not supported by the description in

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

Cochrane, and in fact, is contradicted by the specific and detailed discussion in Cochrane of how materialized views are maintained in response to a change to an underlying base table.

In the 12/27/2006 Office Action, as further purported support for the § 102 rejection over Cochrane, the Examiner cited ¶ [0035] of Cochrane for the proposition that there are two classes of materialized views. As indicated by ¶ [0035], one class of materialized views includes “self-maintainable materialized views” that are maintained incrementally by processing only the set of updated rows without reference to the underlying tables of the materialized view query. The other class of materialized views noted in ¶ [0035] of Cochrane includes materialized views that are maintained incrementally by processing the updated rows and visiting the underlying tables to recompute some portions of the materialized view query.

Based on these teachings of ¶ [0035], the Examiner stated that “the underlying tables have an impact (with respect to updates) to materialized views.” 12/27/2006 Office Action at 23. The fact that underlying base tables “have an impact” on materialized views bears little relevance to the claim elements at issue, namely that the second result set is updated by computing a change to the second result set based on the change to the first result set, and where the change to the first result set is based on the change in the at least one base table.

In fact, the specific teaching of Cochrane is that modifications (in the form of insert/delete/update) on the base tables cause incremental maintenance of summary tables (views) by computing a “raw delta.” Cochrane, ¶ [0093]. The key issue argued by Appellant is that Cochrane does *not* teach the updating of the second result set (for a group-by on a second grouping set) by computing a change to the second result set *based on the change to the first result set (for a group-by on a first grouping set)*, where both the first result set and second result

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

set are contained in the same view that contains results of a cube-based operation on at least one base table. In view of this, it is clear that the anticipation rejection is defective.

The Examiner made the following further erroneous statement in the 12/27/2006 Office Action: "in summary, Cochrane teaches that underlying tables can be visited for recomputation which is the same as computing a change to a second result set based on the change to the first result set." 12/27/2006 Office Action at 23. Computing a change to a result set based on visiting *underlying tables* as taught by Cochrane corresponds to the first update task of claim 1: "update the first result set by computing a change to the first result set based on a change in the at least one base table." The teaching in Cochrane that underlying tables are visited to compute changes to the view does *not* constitute the second update task of claim 1: "update the result set by computing a change to the second result set *based on the change to the first result set*." Note that the second update task differs from the first update task in that the second update task updates the second result set based on the change to the first result set, whereas the first update task updates the first result set based on a change to a base table. Cochrane's teachings regarding computing a change to the view based on the underlying table corresponds only to the first update task of claim 1.

In addition, the "computing deltas of steps I and II of paragraphs [0093] and [0104]" referred to on page 23 of the 12/27/2006 Office Action refer to the delta aggregation noted by Cochrane to produce a "delta cube" according to the complex group-by expression of the view definition. Cochrane, ¶ [0110]. The delta cube has a one-to-one correspondence between the content of the delta cube and the rows in the view to be updated (so that a left outer-join can be performed). Cochrane, ¶ [0114]. In other words, the result sets of the delta cube specifically reflect the changes that are to be made to the view. These result sets in the delta cube are all

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

directly computed from the raw delta stream, which represents changes of the underlying base tables. There is no teaching of one result set being computed based on changes to another result set of the delta cube.

Finally, in the Advisory Action dated March 26, 2007, the Examiner referred to the "hierarchical data cube" referenced in ¶¶ [0064] and [0087] of Cochrane. The fact that Cochrane discloses a hierarchical summary table (cube) does not change the fact that the detailed discussion of incremental maintenance of summary tables (cubes) discussed in ¶¶ [0089]-[0110] of Cochrane, which comes after the reference to the hierarchical data cubes made in ¶¶ [0064] and [0087], is completely different from the update tasks performed in claim 1, as discussed in detail above.

In view of the foregoing, it is respectfully submitted that claim 1, and its dependent claims, are not anticipated by Cochrane.

Independent claims 13 and 22, and their respective dependent claims, are also not anticipated by Cochrane.

In view of the foregoing, reversal of the final rejection of the above claims is respectfully requested.

2. Claims 3, 15, 16, and 24.

Claims 3, 15, 16, and 24 depend from respective independent claims 1, 13, and 22, and are thus allowable for at least the same reasons as those claims. Moreover, claim 3 recites that the view further contains a third result set for a group-by on a third grouping set having a third number of grouping attributes, where the third number is less than a second number of attributes for the second grouping set, and that the controller updates the third result set by computing a change to the third result set *based on the change to the second result set*. As discussed above,

Appn. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

Cochrane performs delta aggregation to produce a delta cube according to the group-by expression of the view definition. Cochrane, ¶ [0110]. The delta cube is then used to perform incremental update of a view, which clearly does not provide the multiple levels of updates recited in claim 3: the first result set updated based on the change to a base table, the second result set updated based on a change to the first result set, and the third result set updated based on a change to the second result set.

Claim 3 is thus further distinguishable over Cochrane for the foregoing reasons. Dependent claims 15, 16, and 24 are further allowable for similar reasons as stated above. In view of the foregoing, reversal of the final rejection of the above claims is respectfully requested.

B. Claims 6-12, 17-21, And 25-28 Rejected Under 35 U.S.C. § 103 Over Cochrane In View Of U.S. Patent No. 7,035,843 (Bellamkonda).

1. Claims 6, 10-12, 17, 18, 25, and 26.

In view of the defective rejection of base claims over Cochrane, it is respectfully submitted that the obviousness rejection of dependent claims 6, 10-12, 17, 18, 25, and 26 over Cochrane and Bellamkonda is also defective.

Therefore, reversal of the final rejection of the above claims is respectfully requested.

2. Claims 7-9.

In view of the defective rejection of base claim 1 over Cochrane, it is respectfully submitted that the obviousness rejection of claim 7 over Cochrane and Bellamkonda is also defective.

Moreover, it is respectfully submitted that a *prima facie* case of obviousness has not been established with respect to claim 7 for at least the additional reason that the hypothetical combination of Cochrane and Bellamkonda does not teach or hint at all elements of claim 7.

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

Specifically, claim 7 recites that the controller is adapted to distribute rows in the first result set across the access modules based on a hash of columns of the second grouping set and at least another column *that is assigned a predefined value*. In the 12/27/2006 Office Action, the Examiner conceded that Cochrane fails to disclose the subject matter of claim 7, but asserted that Bellamkonda discloses the claim features missing from Cochrane. Specifically, the Examiner cited column 17, lines 25-26, of Bellamkonda, which refers to data records that are partitioned on non-rollup keys before being passed from one processing stage to the next processing stage. This teaching of Bellamkonda is distinguished from what is recited in claim 7, where rows of a first result set are distributed across access modules based on a hash of columns of a second grouping set and at least another column *that is assigned a predefined value*. As disclosing the column that is assigned a predefined value, the Examiner referred to column 12, lines 26-28, of Bellamkonda, which states that a single field in the result table at each stage serves as a grouping distinguisher. However, note that this single field is not used for the purpose of performing the hashing distribution mentioned in column 17 of Bellamkonda. In the distribution of results described in column 17 of Bellamkonda, none of the keys that are used for partitioning are assigned a predefined value.

Therefore, it is clear that the hypothetical combination of Cochrane and Bellamkonda does not teach or hint at all elements of claim 7 and its dependent claims. Therefore, the obviousness rejection of claim 7, and its dependent claims, is defective for this additional reason.

Reversal of the final rejection of the above claims is respectfully requested.

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

3. Claims 19-21, 27, and 28.

In view of the defective rejection of base claims 13 and 22 over Cochrane, it respectfully submitted that the obviousness rejection of dependent claims 19-21, 27, and 28 over Cochrane and Bellamkonda is also defective. Moreover, a *prima facie* case of obviousness has not been established with respect to claim 19 and its dependent claim for the additional reason that the hypothetical combination of Cochrane and Bellamkonda does not teach or hint at all elements of claim 19.

Claim 19 recites that the first result set corresponds to a group-by on a first grouping set having N columns, and the second result set corresponds to a group-by on a second grouping set having N-1 columns, and that distributing the first result set to compute the second result set comprises distributing the first result set based on a hash of the N columns, with the column in the first grouping set not present in the second grouping set *being assigned a predefined value*. As discussed above, the distribution performed in column 17 of Bellamkonda clearly does not disclose assigning a predefined value to a column of a grouping set. Therefore, claims 19-21 are allowable over Cochrane and Bellamkonda.

Claims 27 and 28 are allowable for similar reasons.

Reversal of the final rejections of the above claims is respectfully requested.

JUN-25-2007 MON 02:14 PM TROP, PRUNER & HU, PC

FAX NO. 7134688883

P. 16/23

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

RECEIVED
CENTRAL FAX CENTER

JUN 25 2007

CONCLUSION

In view of the foregoing, reversal of all final rejections and allowance of all pending claims is respectfully requested.

Respectfully submitted,

Date: June 26, 2007



Dan C. Hu
Registration No. 40,025
TROP, PRUNER & HU, P.C.
1616 South Voss Road, Suite 750
Houston, TX 77057-2631
Telephone: (713) 468-8880
Facsimile: (713) 468-8883

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

JUN 25 2007

VIII. APPENDIX OF APPEALED CLAIMS

The claims on appeal are:

1. 1. A database system comprising:
 2. a storage to store a view containing results of a cube-based operation on at least one base table, the view containing a first result set for a group-by on a first grouping set, and a second result set for a group-by on a second grouping set; and
 5. a controller, in response to a change to the at least one base table, to:
 6. update the first result set by computing a change to the first result set based on a change in the at least one base table; and
 8. update the second result set by computing a change to the second result set based on the change to the first result set.
 1. 2. The database system of claim 1, wherein the first grouping set has a first number of grouping attributes, and the second grouping set has a second number of grouping attributes, the first number being greater than the second number.
 1. 3. The database system of claim 2, wherein the view further contains a third result set for a group-by on a third grouping set having a third number of grouping attributes, the third number less than the second number,
 4. the controller to further update the third result set by computing a change to the third result set based on the change to the second result set.
 1. 4. The database system of claim 1, wherein the view contains results of a group-by cube operation.
 1. 5. The database system of claim 1, wherein the view contains results of a group-by partial cube operation.

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

- 1 6. The database system of claim 1, further comprising plural access modules and plural storage modules, the access modules to enable parallel access of data in the plural storage modules.
- 1 7. The database system of claim 6, wherein the controller is adapted to distribute rows in the first result set across the access modules based on a hash of columns of the second grouping set and at least another column that is assigned a predefined value.
- 1 8. The database system of claim 7, wherein the view contains results of a cube operation specified by a cube function on plural columns, the at least another column being one of the plural columns of the cube function that is not in the second grouping set.
- 1 9. The database system of claim 7, wherein the view is distributed across the access modules such that plural portions of the view reside in respective storage modules, and wherein the rows in the first result set are distributed across the access modules according to the hash to enable:
 - 5 each access module to locally perform a merge and aggregate operation on the rows of the first result set to produce rows of the second result set; and
 - 7 each access module to locally merge the rows of the second result set into a respective portion of the view without having to first redistribute the rows of the second result set.
- 1 10. The database system of claim 1, wherein the controller is adapted to further:
 - 2 receive a query specifying a group-by operation; and
 - 3 determine whether an answer for the query specifying the group-by operation can be satisfied from the view.
- 1 11. The database system of claim 10, wherein the query specifies a group-by operation on grouping sets S, and the view contains result sets for grouping sets C,
 - 3 the controller to determine whether S is a subset of C to determine whether the answer for the query can be satisfied from the view.

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

1 12. The database system of claim 11, wherein the controller is adapted to modify a WHERE
2 clause of the query in response to determining that S is a subset of C.

1 13. A method for use in a database system, comprising:
2 storing a view containing results of a cube-based operation on at least one base table, the
3 view containing result sets for group-bys on respective grouping sets;
4 updating a first result set by computing a change to the first result set based on a change
5 in the at least one base table; and
6 updating a second result set by computing a change to the second result set based on the
7 change to the first result set.

1 14. The method of claim 13, wherein updating the first result set comprises updating the first
2 result set for the group-by on a first grouping set that has a greater number of columns than a
3 second grouping set corresponding to the second result set.

1 15. The method of claim 13, further comprising updating a third result set by computing a
2 change to the third result set based on the change to the second result set.

1 16. The method of claim 15, further comprising updating a fourth result set by computing a
2 change to the fourth result set based on the change to the third result set.

1 17. The method of claim 13, wherein the database system has plural storage modules to store
2 respective portions of the view, and plural access modules to manage access of respective
3 storage modules,
4 wherein updating the first result set and second result set are performed in parallel by the
5 plural access modules.

1 18. The method of claim 17, further comprising distributing rows of the first and second
2 result sets across the plural access modules.

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

- 1 19. The method of claim 18, wherein the first result set corresponds to a group-by on a first
- 2 grouping set having N columns, and the second result set corresponds to a group-by on a second
- 3 grouping set having N-1 columns, and wherein distributing the first result set to compute the
- 4 second result set comprises distributing the first result set based on a hash of the N columns, with
- 5 the column in the first grouping set not present in the second grouping set being assigned a
- 6 predefined value.
- 1 20. The method of claim 19, further comprising:
 - 2 updating a third result set by computing a change to the third result set based on the
 - 3 change to the second result set, wherein the third result set corresponds to a group-by on a third
 - 4 grouping set having N-2 columns,
 - 5 wherein distributing the second result set across the access modules to compute the third
 - 6 result set is based on a hash of the N columns, with the columns in the first grouping set not
 - 7 appearing in the third grouping set each being assigned to the predefined value.
- 1 21. The method of claim 20, wherein storing the view comprises storing a view for a cube
- 2 operation based on a cube function of the N columns.
- 1 22. An article comprising at least one storage medium containing instructions that when
- 2 executed cause a database system to:
 - 3 store a view containing results of a cube-based operation on at least one base table, the
 - 4 view containing result sets for group-bys on respective grouping sets;
 - 5 update a first result set by computing a change to the first result set based on a change in
 - 6 the at least one base table; and
 - 7 update a second result set by computing a change to the second result set based on the
 - 8 change to the first result set.
- 1 23. The article of claim 22, wherein updating the first result set comprises updating the first
- 2 result set for the group-by on a first grouping set that has a greater number of columns than a
- 3 second grouping set corresponding to the second result set.

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

1 24. The article of claim 22, wherein the instructions when executed cause the database
2 system to further update a third result set by computing a change to the third result set based on
3 the change to the second result set.

1 25. The article of claim 22, wherein the database system has plural storage modules to store
2 respective portions of the view, and plural access modules to manage access of respective
3 storage modules,

4 wherein updating the first result set and second result set are performed in parallel by the
5 plural access modules.

1 26. The article of claim 25, wherein the instructions when executed cause the database
2 system to further distribute rows of the first and second result sets across the plural access
3 modules.

1 27. The article of claim 26, wherein the first result set corresponds to a group-by on a first
2 grouping set having N columns, and the second result set corresponds to a group-by on a second
3 grouping set having N-1 columns, and wherein distributing the first result set to compute the
4 second result set comprises distributing the first result set based on a hash of the N columns, with
5 the column in the first grouping set not present in the second grouping set being assigned a
6 predefined value.

1 28. The article of claim 27, wherein storing the view comprises storing a view for a cube
2 operation based on a cube function of the N columns.

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

IX. EVIDENCE APPENDIX

None.

Appln. Serial No. 10/706,656
Appeal Brief Under 37 C.F.R. § 41.37

X. RELATED PROCEEDINGS APPENDIX

None.